arXiv:1402.0170v1 [cs.CV] 2 Feb 2014

Collaborative Receptive Field Learning

Shu Kong AIMERYKONG@GMAIL.COM
HKUST, and Noah's Ark Lab of Huawei Co. Ltd.

Zhuolin Jiang ZHUOLIN.JIANG@HUAWEI.COM
Qiang Yang QIANG.YANG@HUAWEI.COM
Noah's Ark Lab of Huawei Co. Ltd.

Abstract

The challenge of object categorization in images is largely due to arbitrary translations and scales of the foreground objects. To attack this difficulty, we propose a new approach called collaborative receptive field learning to extract specific receptive fields (RF's) or regions from multiple images, and the selected RF's are supposed to focus on the foreground objects of a common category. To this end, we solve the problem by maximizing a submodular function over a similarity graph constructed by a pool of RF candidates. However, measuring pairwise distance of RF's for building the similarity graph is a non-trivial problem. Hence, we introduce a similarity metric called pyramid-error distance (PED) to measure their pairwise distances through summing up pyramid-like matching errors over a set of low-level features. Besides, in consistent with the proposed PED, we construct a simple non-parametric classifier for classification. Experimental results show that our method effectively discovers the foreground objects in images, and improves classification performance.

1. Introduction

It is widely known that the difficulty in automatic object categorization from images is largely due to the arbitrary translations and scales of the foreground objects. To solve the problem, researchers have designed robust image features like SIFT (Lowe, 2004) and HoG (Dalal & Triggs, 2005) for image representation, reliable image (region) matching techniques such as image alignment (Kim & Xing, 2013) and detection (Russakovsky et al., 2012), and sophisticated classifiers (Duchenne et al., 2011).

Good image representation first concerns robust features. Current feature learning methods propose to learn mid-level features hierarchically built over low-level ones, which are also preferably learned adaptively rather than the hand-crafted ones (Girshick et al., 2013), *e.g.* SIFT and HoG. Then, various representation learning methods are proposed, such as spatial pyramid (Lazebnik et al., 2006) and multiple layers of pooling and downsampling (Krizhevsky et al., 2012). These representations can roughly preserve the salient object structures, thus they enhance the discriminativeness of image representations. With the adaptively learned features and discriminative representations, the classification performance is improved accordingly. But as these methods cannot effectively handle large translations and scales of the objects, the accuracy gains are still limited.

There are some approaches attempting to localize the foreground objects for better encoding images, such as saliency detection (van de Sande et al., 2011; Fu et al., 2013), segmentation (Chang et al., 2011) and object detection (Nguyen et al., 2009; Russakovsky et al., 2012). Essentially, these methods can be cast as the so-called receptive field learning as they intend to find the most desirable image regions (receptive fields) for particular tasks. For example, Jia *et al.* try to solve this problem by optimizing spatial pyramid matching (SPM) in building mid-level features (2012). Their method selectively combines pooled features in predefined image regions to improve the discriminability of overall image representations. However, Their method using mid-level features learns the same combination patterns for all images of different categories, thus it still fails in handling prominent translations and scales in individual images.

To effectively handle arbitrary scales and translations of objects, we propose a new framework called **collaborative receptive fields learning** (coRFL), which intends to discover specific receptive fields (RF's) or image regions that

The work is done when Shu Kong is jointly with HKUST and Noah's Ark Lab of Huawei Co. Ltd.

mainly cover the foreground objects from the same category. coRFL merely requires weak labels (Russakovsky et al., 2012); that is to say, there is no exact object location information but the category-level label for each image. Moreover, coRFL learns to find these RF's collaboratively among multiple images from a common category, thus leading to reciprocal accuracies for discovering their common object. Note that some definitions of receptive field in neural science are different to ours (Olshausen et al., 1996), but we keep using this term to highlight the meaning that RF's in images received by the computer should capture the most distinct foreground object. We model coRFL as selecting specific vertices from a graph, which is constructed by pairwise similarities of RF candidates from these images. Borrowing some vision-based priors, we formalize the problem as a submodular function, with which a simple greedy method suffices to produce performance-guaranteed solutions.

However, in building the graph, finding the right metric of pairwise distance between RF's with varying sizes is a nontrivial problem. One intuitive way to represent the RF is to use the mid-level feature concatenated by multi-layer pooled vectors (Yang et al., 2009) with the same length. But these features usually have thousands of dimensions, and will lose distinct information due to its vector quantization or sparse coding (Domingos, 2012; Boiman et al., 2008). For this reason, we introduce Pyramid-Error Distance (PED), a nonparametric method to measure the distance of image regions over sets of low-level SIFT features.

We perform coRFL in training images of each category to purify the training set by only preserving RF's that capture the meaningful foreground objects. With the proposed PED, we design a nonparametric classifier to match qeury images with the purified training set. Through experiments over both synthetic data and benchmark databases, we show the effectiveness of our proposed framework.

Contributions and Paper Organization: We first review essential preliminaries in Section 2. Then we elaborate our framework of coRFL in Section 3, the metric of Pyramid-Error Distance (PED) in Section 4, and our designed nonparametric classifier in Section 5, respectively. We evaluate our framework with experiments in Section 6, before concluding in Section 7.

2. Related Work

There are three keywords in our framework, receptive field learning, submodular function and similarity metric.

Receptive Field Learning: Multiple problems in computer vision can be seen as receptive field learning to aid image understanding. For example, saliency detection aims to discover regions that capture human attention in the images with perceptual biases; the result of detected salient regions anticipate better image matching by only considering these salient regions (Fu et al., 2013). Besides, image segmentation aim to simplify or change the representation of an image into something that is more meaningful and easier to analyze (Shi & Malik, 2000). Requiring the localization information, van de Sande *et al.* use salient regions through multi-scale segmentation with multiple cues to detect the object of interest (2011). Moreover, by considering object translations and scales, Russakovsky *et al.* propose object-centric spatial pooling (OCP) approach (2012), which first infers the location of the objects and then uses their locations to pool foreground and background regions separately to form the mid-level features. OCP learns the object detectors with the weak labels, *i.e.* there is no exact object location information in images. This is the same condition in our work.

In particular, Jia *et al.* explicitly work on receptive field learning through learning to selectively combine pooled vectors over 100 predefined grids (2012), as demonstrated by Fig. 1 (b). But their method learns the same combination pattern across all images from different categories. Therefore, when facing notable translation and scale changes of foreground objects, it cannot be guaranteed to achieve improved performance. Moreover, Duchenne *et al.* introduce a graph-matching kernel (GMK) to address object deformations between every pair of images (2011). However, the kernel requires more time to calibrate images for large-scale datasets, and also fails in calibrating images with extremely cluttered backgrounds

Submodular Function: The natural and wide applicability of submodular function makes it receiving more attention in recent years (Iyer et al., 2013). Let $\mathcal{V}$ be a finite ground set. A set function $\mathcal{F} : 2^{\mathcal{V}} \rightarrow \mathbb{R}$ is submodular if $\mathcal{F}(A \cup a) - \mathcal{F}(A) \geq \mathcal{F}(A \cup \{a, b\}) - \mathcal{F}(A \cup b)$, for all $A \subseteq \mathcal{V}$ and $a, b \in \bar{A}$. Here, $\bar{A} = \mathcal{V}/A$ is the complement of A. The property is referred to as diminishing return property, stating that adding an element to a smaller set helps more than adding it to a larger one. As for other properties of submodularity, please refer to (Iyer et al., 2013) and references therein.

Similarity Metric: The similarity measurement of image regions is still an open problem. Recently, the combination of mid-level features and SVM classifier generally produce promising results (Lee et al., 2009; Boureau et al., 2010; Zeiler et al., 2011). The mid-level features are usually generated by concatenating multi-layer (pooled) features within spatial pyramid pattern as demonstrated in Fig. 1 (a) (Lazebnik et al., 2006; Yang et al., 2009; Coates & Ng, 2011), or learned through the convolutional neural networks (CNN) (Krizhevsky et al., 2012; Girshick et al., 2013). However, pairwise similarity between mid-level

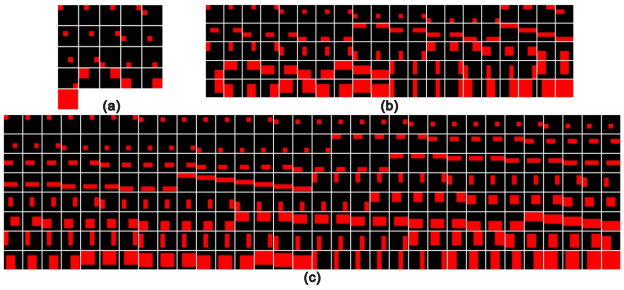

Figure 1. Comparisons of receptive field (RF) candidates: (a) spatial pyramid grids (Yang et al., 2009); (b) RF's for selective combination (Jia et al., 2012); (c) RF candidates in our method.

features cannot be reliably measured by their Euclidean distance, due to both their high dimension and the vector quantization or sparse coding stage in extracting mid-level features (Boiman et al., 2008; Domingos, 2012). Among these methods, the Euclidean distance over low-level SIFT descriptors (Boiman et al., 2008) motivates our proposed metric.

3. Collaborative Receptive Field Learning

In this section, we present the proposed framework of collaborative receptive field learning (coRFL) in detail. With the weak labels and fed multiple images of a common category, coRFL collaboratively extracts specific receptive field (RF's) that capture the common foreground objects. Solving the problem is the core in our proposed framework, because we perform coRFL over training images in each category to purify the training set for matching queries, and the resultant images only preserve the most meaningful foreground objects. We first demonstrate how to extract RF candidates, and then present some vision-based priors before the formalism of coRFL.

3.1. Extracting Receptive Field Candidates

Suppose there are N images available from a specific category, without loss of generality, we predefine m templates to extract RF candidates in images. In this work, we define $m = 256$ candidates as shown in Fig. 1 (c), leading to $M = mN$ candidates in total for these N images. In contrast to the approach that defines 100 grids (Jia et al., 2012), our overlapping grids can capture the foreground objects more reliably and correctly, thus preventing the computation from covering only object parts or too many image backgrounds. Now, we solve coRFL by selecting the most desirable image regions or receptive fields (denoted by RF_+) that capture the common and distinct objects from these images. In particular, we can specify the number of selected RF_+'s as K. Then, the crucial point is to sketch a mechanism to find the most desirable RF_+'s that are distinct from the negative ones (denoted by RF_- that mainly covers the cluttered background).

It is worth noting the difference between our RF candidates and the pooled features of image grids (Lazebnik et al., 2006). Specifically, our method finds the most desired RF's that capture the foreground objects from predefined grids; and these selected RF's form the training set used for classification. In contrast, SPM-based methods (selectively) concatenate the pooled vectors for the final representation of the overall image, such as in (Yang et al., 2009; Jia et al., 2012). As a result, only our method explicitly considers object scale and translation in individual images.

In addition, Girshick *et al.* propose to extract region proposals (with different sizes) in images for object detection (2013). These region proposals can also be seen as RF candidates, but they are required to be warped with brutal force into a fixed size, so that they can be fed into a CNN (Krizhevsky et al., 2012). Different from these methods, this transformation may destroy information related to object appearance and shape. In contrast, our method preserves such valuable information by allowing various sizes of the RF's.

3.2. Inter- and Intra-Image Prior

Inspired by the area of saliency detection, we assume that the RF_+ capturing the object is more salient than others (RF_-'s), which mainly cover the background or object parts. In other words, there should be a large contrast between RF_+ and RF_-, both from the image itself and other images. We call intra- and inter-image prior, respectively. Therefore, an oracle should find (a few) RF_+'s which have small similarities with (most) RF_-'s, *i.e.* pairwise similarities between selected RF_+ and RF_- should be small.

Besides, as multiple images are given from a common category, we can say each image has at least one RF_+ that makes them correlated semantically by capturing the common objects. Based on this repeatedness principle, inter-image relationship can be exploited by considering similarities between RF_+'s from different images should be large. Moreover, pairwise similarities between RF_-'s should be small, since RF_-'s mainly capture cluttered background which can be seen as noises.

To model the priors, we build a graph $\mathbf{S} \in \mathbb{R}^{M \times M}$ to record the similarities between each pair of RF candidates. The larger element, say s_{ij} in $\mathbf{S}$, means that receptive fields i and j are more similar to each other. As the measurement is a nontrivial problem over RF's, we continue to elaborate coRFL and put graph construction in Section 4. Now, an oracle will find a set of vertices as the RF_+'s indexed by A, such that the sum of similarities within RF_+'s is a maxima. Meanwhile, the summed similarity between

RF_+'s and RF_-'s indexed by the complement $\bar{A}$, as well as summed similarity within RF_-'s, is a minima.

We define the following operation over matrix $\mathbf{S}$ with two sets A and B indexing rows and columns, respectively:

$$\mathbf{S}_{A,B} = [s_{ij}] \in \mathbb{R}^{|A|\times|B|}, \forall i \in A, j \in B. \tag{1}$$

Note that $\mathbf{S}_{A,B} = \mathbf{S}_{B,A}^T$ as we require the symmetric matrix $\mathbf{S}$ to specify an undirected graph. Moreover, over $\mathbf{S}_{A,B} \in \mathbb{R}^{|A|\times|B|}$, we define a function $h(\mathbf{S}_{A,B})$ for the sum of pairwise similarities between A and B as:

$$h(\mathbf{S}_{A,B}) = \begin{cases} \sum_{i\in A}\sum_{j\in B} c_{ij}, & A \neq \varnothing \text{ and } B \neq \varnothing, \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

Therefore, we can say the RF_+'s indexed by A should simultaneously lead to a maxima of $h(\mathbf{S}_{A,A})$, a minima of $h(\mathbf{S}_{A,\bar{A}})$ and $h(\mathbf{S}_{\bar{A},\bar{A}})$. By unifying all these terms, an oracle should find A that leads to the maxima of the following:

$$\mathcal{H}(A) = \log\left(\mu + h(\mathbf{S}_{A,A}) - \alpha h(\mathbf{S}_{A,\bar{A}}) - \beta h(\mathbf{S}_{\bar{A},\bar{A}})\right), \tag{3}$$

in which μ is a constant scalar[1] that is sufficient large to ensure $\mu + h(\mathbf{S}_{A,A}) - \alpha h(\mathbf{S}_{A,\bar{A}}) - \beta h(\mathbf{S}_{\bar{A},\bar{A}})$ is positive; and positive parameters α and β jointly control relative importance of the terms. The $\log$ operation in Eq. 3 is used to ensure the following submodular property[2]:

Proposition 1 *There exist proper α and β that make the proposed function $\mathcal{H}(A)$ in Eq. 3: $2^{\mathcal{V}} \to \mathbb{R}$ a monotonically increasing and submodular function.*

Especially, we have:

Lemma 1 *For any $\tau \in \mathbb{R}^*$, $\mathcal{H}(A)$ in Eq. 3 is a monotonically increasing and submodular function by setting $\alpha = \tau - 1$ and $\beta = \tau$.*

Following Lemma 1, we set $\tau > 1$ and require $\beta = \tau$ and $\alpha = \tau - 1$ to benefit from the desirable submodularity and monotonicity of $\mathcal{H}(A)$, and to model the principles in finding RF_+'s. Hereafter, we rewrite as $\mathcal{H}_\tau(A)$ to explicitly highlight the sole parameter τ.

3.3. Balance Penalty

Since the common weak labels enable images from one category to be correlated to each other, it is reasonable for each image to contribute one RF_+ by itself. Thus, we should extract at least one RF_+ from each training image. A benefit of this balance is the preservation of intra-class variability, which helps alleviate the overfitting problem.

We call this balance principle, *i.e.* the number of RF_+ in each image should be balanced. Specifically, let A_j index the RF_+'s from the j^{th} image, and $A = \cup_{j=1}^N A_j$ indices all the positive RF_+'s. We add a penalty term to our objective function as below to balance the number of RF_+'s in the images:

$$\begin{aligned} \mathcal{G}(A) &= \sum_{j=1}^{N} \log(|A_j| + 1), \\ \text{s.t. } &\cup_{j=1}^N A_j = A, A_i \cap A_j = \varnothing, \forall i \neq j, \end{aligned} \tag{4}$$

where $|A_j|$ means the cardinality of index set A_j. Particularly, we have $|\varnothing| = 0$. To understand the functionality of Eq. 4, please consider the following proposition:

Proposition 2 *With the function defined as below over vector $\mathbf{x} = [x_1, \ldots, x_N] \succcurlyeq 0$:*

$$g(\mathbf{x}) = \textstyle\sum_{j=1}^N \log(x_j + 1),$$

if $x_c \leq x_i, \forall i \neq c$, then $g(x_c + 1|\mathbf{x}) \geq g(x_i + 1|\mathbf{x})$. where $g(x_i + 1|\mathbf{x}) = \sum_{j\neq i} \log(x_j + 1) + \log(x_i + 1 + 1)$.

This property demonstrates that adding smaller elements achieves greater reward. In particular, over the graph defined by $\mathbf{S}$, the vertices (RF's) are preferred to be selected from images one after another. As a result, balancing the number of elements can be achieved.

Furthermore, we have the following proposition:

Proposition 3 *The function $\mathcal{G}(A)$ in Eq. 4: $2^{\mathcal{V}} \to \mathbb{R}$ is a monotonically increasing and submodular function.*

3.4. Center-Bias Principle

Inspired by the saliency detection, we exploit center-bias principle (Tatler, 2007) to mildly constrain the RF_+ w.r.t its location in the image. Specifically, an RF_+ appears around the center of the image with high probability. Our mild center-bias constraint means that, searching for RF_+'s should focus more around the image center, but still allows to capture the most desired RF_+ locating near the image margin with high fidelity.

Intuitively, center bias can be modeled through the position of RF's. Let $\boldsymbol{\theta} \in \mathbb{R}^M$ denote the distances[3] of all the RF's to the image center, specifying center-bias constraint for each of the M RF candidates. One intuitive example is to constrain that $\|\boldsymbol{\theta}_A\|_1$ to be small, where $\boldsymbol{\theta}_A$ means a sub-vector comprising of elements indexed by A. We specify

[1] Generally, we set $\mu = 1 + \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}})$, where τ is defined as in Lemma 1.

[2] All proofs to propositions and lemmas are presented in the supplementary material.

[3] Here the distance does not necessarily mean the Euclidean distance. For proper constraint, Euclidean distance with a Gaussian kernel is preferred in our work.

$\|\mathbf{q}_A\|_1 = 0$ for $A = \varnothing$. Alternatively, if we store the reciprocal of center distances in $\mathbf{q} = [q_k] = [\frac{1}{\theta_k}], k = 1, \ldots, M$, then we need to constrain that A leads to a relatively larger $\|\mathbf{q}_A\|_1$, maximizing which pushes our mechanism to focus on RF's around the images' centers.

3.5. Objective Function

With the proposed inter- and intra-image prior, the balance penalty $\mathcal{G}(A)$, and the mild center-bias penalty, we turn to maximize the following objective function to find RF_+'s indexed by A:

$$A = \underset{A \in \mathcal{I}}{\operatorname{argmax}} \left\{ \mathcal{F}(A) \equiv \mathcal{H}_\tau(A) + \lambda_1 \mathcal{G}(A) + \lambda_2 \|\mathbf{q}_A\|_1 \right\} \quad \text{s.t. } |A| \leq K, \tag{5}$$

where λ_1 and λ_2 are the parameters to control inter- and intra-image prior term, the balance penalty term and center prior term, respectively. With the proposition as below, we can see that exactly K receptive fields are extracted.

Proposition 4 *The function $\mathcal{F}(A)$ in Eq. 5: $2^{\mathcal{V}} \rightarrow \mathbb{R}$ is a monotonically increasing and submodular function, and induces a uniform matroid $\mathcal{M} = (\mathcal{V}, \mathcal{I})$, where $\mathcal{V}$ is the point set, and $\mathcal{I}$ is the collection of subsets $A \subseteq \mathcal{V}$.*

Submodularity described by the above proposition indicates that a simple greedy algorithm suffices to produce performance-guaranteed solutions with a theoretical approximation $(1 - 1/e)$ (Nemhauser et al., 1978). The greedy search requires $|\mathcal{V} - A|$ evaluations for the marginal gains before adding a new element into A at each iteration. To speed up the optimization process, we use the lazy greedy (Leskovec et al., 2007) by constructing a heap structure over marginal gains of each element. Even through the addition of any element into A impacts the gains of the remaining ones, we can merely update the gain of the top element in the heap, instead of recomputing the gains for every remaining element. The key idea is that the gain for each element can never increase due to the diminishing return property of submodular function, which can be illustrated by the naive search method in Fig. 3. Moreover, the recomputation of gain for the top element in the heap is not much smaller in many cases, hence the top element will stay the top element even after the update.

The worst case is to update the gain for each element and then re-establish the heap after the addition of any elements to A, leading to the complexity $O(|\mathcal{V}| \log |\mathcal{V}|)$ for rebuilding the heap, and the overall complexity $O(|\mathcal{V}|^2 \log |\mathcal{V}|)$ (Cormen et al., 2009) of the optimization. But in practice, the lazy algorithm only requires a few updates in the heap at each iteration. Hence, the complexity of the optimization is effectively $O(|\mathcal{V}| \log |\mathcal{V}|)$.

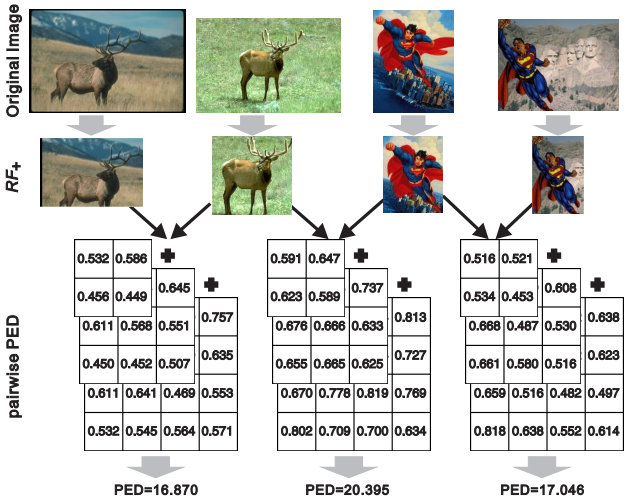

Figure 2. Illustration of Pyramid-Error Distance (PED). The first row presents the original images, and their corresponding RF's are learned by our algorithm in the second row. The third row shows the pairwise PED. PED encourages the similarity between RF's from the same class is larger than that from different classes.

4. Similarity Graph Construction via Pyramid-Error Distance

In this section, we investigate how to measure the pairwise similarity of RF's in constructing the graph. As the RF's are essentially image regions with varying sizes, measuring them is a nontrivial problem. One intuitive idea is to borrow the mid-level pooled features (Jia et al., 2012) to represent each RF candidates, as the pooling process generates feature vectors with fixed length. But it will produce disastrous results due to both high dimensionality and vector quantization or sparse coding. Therefore, we introduce a new metric called Pyramid-Error Distance (PED).

4.1. Pyramid-Error Distance

We split each RF candidate into pre-defined grids at multiple levels. For instance, we use three partition scale 2×2, 3×3 and 4×4, leading to $L = 29$ grids in total. Please note that the pyramid partitions is done in each single RF, instead of the overall image. This is totally different from SPM-based methods (Lazebnik et al., 2006; Boureau et al., 2010; Jia et al., 2012) that concatenate the pooled vectors of all grids to represent the whole.

Throughout our work, we extracted the low-level SIFT features over each grid of RF. Before measuring the similarity of two RF's, we first calculate the distance of a pair of grids from two RF's at a corresponding position indexed by l. Let $X_l = \{\mathbf{x}_i \in \mathbb{R}^p | i = 1, \ldots, r\}$ and $Y_l = \{\mathbf{y}_j \in \mathbb{R}^p | j = 1, \ldots, q\}$ be two sets consisting of p-dimensional descriptors, representing the two corre-

sponding grids, respectively. As a result, even with various sizes, grids can be represented by sets of low-level features. Please also note the descriptor number r and q are not necessarily equal, due to the number of feature points automatically detected in the RF's of different sizes. We define the distance at set level as below:

$$dist(X_l||Y_l) = \frac{1}{2r}\sum_{i=1}^{r}\big(\min_j \|\mathbf{x}_i - \mathbf{y}_j\|_F^2\big) + \frac{1}{2q}\sum_{j=1}^{q}\big(\min_i \|\mathbf{x}_i - \mathbf{y}_j\|_F^2\big). \tag{6}$$

Furthermore, let $RF_i = \{X_l | l = 1, \ldots, L\}$ and $RF_j = \{Y_l | l = 1, \ldots, L\}$ be two RF's, as shown by the second row in Fig. 2. Then, with the pyramid partitions, we now arrive at the PED between two RF's as:

$$D(RF_i||RF_j) = \sum_l dist(X_l||Y_l), \tag{7}$$

in which l indexes the grid in specific location within a defined pyramid partition. From the third row in Fig. 2, the PED is calculated by the sum of grid distances in a pair of RF's. To analyze the complexity, we naively assume there are n_{grid} descriptors (with d-dimensionality) in each grid, then calculating PED is of complexity $O(dLn_{grid}^2)$.

4.2. Similarity Graph Construction

With the defined PED in Eq. 7, we calculate the similarities of each pair of RF's, and construct the graph $\mathbf{S}$ accordingly. In detail, over two receptive fields RF_i and RF_j, we calculate their PED as $D(RF_i||RF_j)$, and then transform PED into similarity s_{ij} by a Gaussian kernel:

$$s_{ij} = \exp\big(-\frac{D(RF_i||RF_j)}{2\sigma^2}\big), \tag{8}$$

in which σ is the parameter controlling the transformation.

Actually, the similarity graph can be seen as the derivation of a distance graph through the Gaussian kernel. As the distance graph is built by every pair of RF candidates, it is a dense one that connects many uncorrelated candidates. Therefore, to purify the similarity graph, we can either keep a fixed number of smallest values in each row/column of the distance graph, or set a threshold to remove larger values, leading to the so-called kNN graph and ϵ-ball graph (Belkin & Niyogi, 2003), respectively.

It is also worth noting that building the similarity graph is the most costly stage in our computation. The popular methods usually adopt the dense feature extraction scheme (Lazebnik et al., 2006), which, supposedly n descriptors (with d-dimensionality) being extracted in each of the N images from a specific category, requires computational complexity $O(dn^2N^2)$ for constructing the graph based on PED. As the dense extraction of SIFT descriptors consistently leads to thousands of descriptors, constructing the graph is extremely time-consuming. To expedite this stage, we can either turn to fast approximate kNN graph construction (Chen et al., 2009) or the original SIFT feature (Lowe, 2004), which incorporates interest point detection and feature descriptor extraction. But we merely use the original SIFT feature. Essentially, with the interest point detection technique in SIFT, only $n \approx 150$ descriptors are generated in an image of 150×150-pixel resolution. Then, it is efficient enough for calculating pairwise PED among RF candidates and constructing the similarity graph in our experiments. Moreover, in contrast to the dense extraction scheme that produce most unnecessary descriptors, such detection technique leads to more meaningful SIFT descriptors in informative regions.

5. Classifier Design

Our framework is similar to multi-instance learning (Dietterich et al., 1997), but is of particularity, which is especially reflected from principles like center-bias, intra- and inter-image contrast. By solving the problem of collaborative receptive field learning, we design a nonparametric classifier by incorporating RF-to-class metric and center-bias penalty.

With the learned RF's in training set, we put SIFT features of the grids at corresponding positions in a set of pools, and denote P_l^c to store the descriptors from all training images of the c^{th} class at the specific grid indexed by l. Then, fed a query image, our method first extracts SIFT features and M RF candidates (denoted by $\{RF_k\}, k = 1, \ldots, M$, and $RF_k = \{X_l^k | l = 1, \ldots, L\}$). Then, it predicts the label by comparing RF-to-class distances of all the C categories:

$$c^* = \underset{c}{\operatorname{argmin}}\big\{\min_k \sum_{l=1}^{L} dist(X_l^k||P_l^c) + \lambda_2 q_k\big\}. \tag{9}$$

Inspired by (Boiman et al., 2008), we exploit KD-tree (Bentley, 1975) to speed up the classification process, which requires low complexity $O(Nn\log(Nn))$ for training the KD-tree for each category, and has $O(Cn\log(Nn))$ complexity to predict a query image.

6. Experimental Validation

In this section, we first qualitatively validate the effectiveness of our method for coRFL over a synthetic dataset in discovering the RF_+'s from images. Then, we use public benchmarks to quantitatively evaluate our method in object categorization, including Caltech101 (Fei-Fei et al., 2007) and Caltech256 (Griffin et al., 2007). Finally, we discuss the parameters used in our experiments.

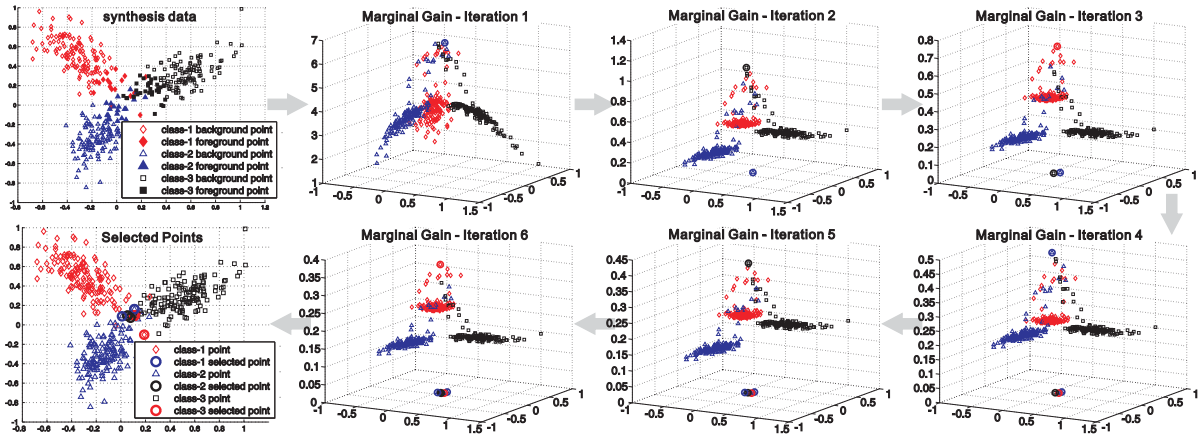

Figure 3. (Best seen in color and zoom in.) Demonstration of the proposed submodular function in exemplar selection over a synthetic data set. Leftmost column: display of the generated data points (upper), and the six selected exemplars by the proposed submodular function (down). The rest columns: marginal gain at each iteration. This figure illustrates that, with our objective function, the most similar or correlated exemplars can be found.

6.1. Synthetic Data

To generate the dataset, we use three Gaussian distributions to produce random points in a 2D plate, as demonstrated by the first upper-left panel in Fig. 3. The three clusters can be seen as three images, and their intersection can be seen as the common objects in the images, meanwhile, points far away from the intersection can be imagined as cluttered backgrounds. Therefore, our objective function in Eq. 5 is expected to find a set of points located in the intersection. Please note that we set $\lambda_2 = 0$ as center-bias prior does not apply to the synthetic data; and set $\lambda_1 = 2$, $\tau = 2$. We use the Euclidean distance of their locations in the 2D plate to build the similarity graph with the Gaussian kernel controlled by parameter $\sigma = 0.3$.

To better understand the process[4], we plot the marginal gains at the first six iterations in Fig. 3, and the six most desirable points in the bottom-left panel. The optimization is done by the native greedy method. From the marginal gain at each iteration, we can see points near the intersection have larger expected gains. This is owing to our inter-image prior in Eq. 3. This figure demonstrates the effectiveness of our approach in finding the most correlated RF's that cover the common foreground object.

6.2. Benchmark Databases

Caltech101 and Caltech256 contain 102 and 256 categories, and have $9,144$ and $30,607$ images, respectively. Caltech256 have higher intra-class variability and higher object location variability (translations and scales) than Caltech101. We resize every image into no more than 150×150-pixel resolution with original aspect ratio. We follow the common setup on the two benchmarks (Yang et al., 2009), *i.e.* 30 images per class are randomly selected as the training set and the rest for testing. We perform coRFL in each category and set $K = 30$ so that exactly 30 RF_+'s are extracted for each class. The average performance after 10 random splits is reported.

We compare our method with several state-of-the-art ones. Most approaches learn mid-level features over low-level ones to represent the overall image, including Convolutional Deep Belief Networks (CDBN) (Lee et al., 2009), adaptive Deconvolutional Networks (DN) (Zeiler et al., 2011), LC-KSVD (Jiang et al., 2013), Kernel SPM

Table 1. Classification accuracies (%) by different methods on the Caltech101 and Caltech256.

Method	Caltech101	Caltech256
CDBN (Lee et al., 2009)	65.4	-
DN (Zeiler et al., 2011)	71.0	33.2 ± 0.8
LC-KSVD (Jiang et al., 2013)	73.6	34.3
KSPM (Lazebnik et al., 2006)	64.6	29.5 ± 0.5
ScSPM (Yang et al., 2009)	73.2	34.0 ± 0.6
LLC (Wang et al., 2010)	73.4	41.2
RFL (Jia et al., 2012)	75.3 ± 0.7	-
GMK (Duchenne et al., 2011)	80.3 ± 1.2	38.1 ± 0.6
NBNN (Boiman et al., 2008)	70.4	37.0
Ours	$\mathbf{83.4 \pm 1.3}$	$\mathbf{45.7 \pm 1.1}$

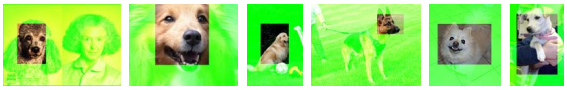

Figure 4. Samples of selected receptive fields over "dog" category from Caltech256 (best seen in color and zoom in).

[4]Code is available at Shu Kong's GitHub: https://github.com/aimerykong/coRFL

(KSPM) (Lazebnik et al., 2006), Sparse Coding based SPM (ScSPM) (Yang et al., 2009), Locality-constrained Linear Coding (LLC) (Wang et al., 2010), Receptive Field Learning (RFL) (Jia et al., 2012), and GMK (Duchenne et al., 2011). CDBN and DN belong to the deep feature learning framework which hierarchically learns adaptive features for image. Both of them generate mid-level features with the spatial pyramid partition and kernel SVM for classification. LC-KSVD is a deeper approach that simultaneously learns a linear classifier and a higher-level dictionary over the mid-level features. The rest methods learn a codebook (consisting of approximate 1024/2048 words for the two benchmarks) over hand-craft descriptors like SIFT, HoG and Macrofeature (Boureau et al., 2010); encode them over the codebook by vector quantization or sparse coding; and then adopt the pooling technique to obtain the feature vectors for image regions w.r.t a 3-layer-pyramid partition. Finally, they concatenate the pooled vectors into a larger one as the image feature and feed into a linear SVM. Additionally, the Naive-Bayes Nearest-Neighbor (NBNN) (Boiman et al., 2008) is closely related to ours, as it directly uses dense SIFT descriptors, image-to-class metric and NN for classification. We list the comparisons in Table 1. An illustration of the selected RF_+'s by our method are displayed in Fig. 4 and the first row in Fig. 5.

As can be seen from Table 1, our method outperforms all the others. It is worth noting the improvement of our method over NBNN attributes to our PED and SIFT extraction with interest point detection. Because PED explicitly considers the shape/structure of the objects and interest point detection removes noisy descriptors. In contrast, NBNN merely constrains position distances of the dense SIFT to be small, thus it incorporates noisy descriptors and fails to handle notable changes of object translation and scale. It is also worth noting that the performance gain brought by our method for Caltech256 is higher than that for Caltech101. We assume the reason is that the changes of object translation and scale is larger in Caltech256 than those in Caltech101. Moreover, we observe that the more cluttered background in the images is, the better performance of our method achieves in finding the objects. This is due to the functionality of our objective function that intends to find the most desirable RF_+'s, and leaves behind the assumed RF_-'s which hold smaller sum of pairwise similarities.

6.3. Parameter Discussion

Our work involves several parameters, including τ, λ_1 and λ_2 in the objective function Eq. 5, and k and σ in the Gaussian kernel for constructing the similarity graph.

τ should be greater than 1 to ensure the physical meaning of our model. When varying the value of τ, we find the classification performance and the visualization of the learned receptive fields do not suffer at all. The reason we guess is due to our objective function, which constrains the assumed RF_-'s have large PED (small similarities). Therefore, larger τ will indirectly contribute to discovering the RF_+'s by finding the most dissimilar RF_-'s. Moreover, a larger value in λ_1 guarantees that each training image contribute at least one RF_+ (overfitting problem is thus alleviated), so we merely set $\lambda_1 = 100$ to ensure this. λ_2 controls the mild center-bias constraint, and has an impact on the performance. Since most foreground objects appear near the center of images, a suitable λ_2 helps to find the real RF_+'s. This can be demonstrated by Fig. 5.

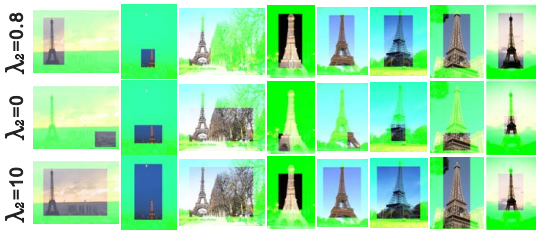

Figure 5. Larger λ_2 means locating the RF_+'s at the center of Caltech256 images with brutal force, while smaller value produces RF_+'s that merely capture object parts. The real RF_+'s can be found with a suitable σ, as the whole object appears in the image center for most images.

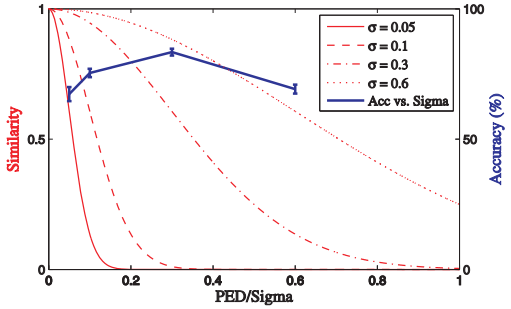

Figure 6. Discussion of σ vs. similarity/accuracy in transforming the dissimilarity graph into the similarity graph.

To construct the similarity graph, we essentially normalize the PED-based distance graph by dividing its largest element, so that all the entries have values in the range of $[0, 1]$. Then, we transform it into similarity graph with Gaussian kernel controlled by σ. We plot the impacts of σ in Fig. 6 over Caltech101 database. Intuitively, with the normalized dissimilarity graph, we can anticipate meaningful outcomes by setting $\sigma \in (0, 1)$. The curve of accuracy vs. σ also demonstrates this intuition. Therefore, we set $\sigma = 0.3$ throughout our work.

7. Conclusion with Discussion

In this work, we introduce a new problem called collaborative receptive field learning (coRFL), which intends to find receptive fields (RF's) or image regions from multiple images that cover the foreground objects of a common category. coRFL merely exploits the weak labels without any exact locations of the objects in the image. By modeling the problem as selecting specific vertices from a similarity graph with consideration of some vision-based priors, we solve coRFL by a submodular function, with which a simple greedy algorithm suffices to produce performance-guaranteed solutions. Furthermore, we propose the Pyramid Error Distance (PED) to measure pairwise distance of RF's. We perform coRFL over images of each category to purify the training set, so that the purified set merely preserves the most meaningful foreground objects. Moreover, in consistent with the PED, we design a simple nonparametric classifier for the final classification.

Our work by no means tries to compete the best results in literature (Zeiler & Fergus, 2013), but it presents several worthwhile research directions.

- In building the similarity graph, we exploit the SIFT feature within the proposed PED. Even though fast graph construction techniques can be explored, other sophisticated representations for the receptive fields are solicited with the consideration of more robustness and efficiency. Especially, learning adaptive features within deep architecture can be exploited to represent the image regions (Kong et al., 2014), and learning adaptive metric for matching is also a research direction for image region matching.
- We model the problem with a simple submodular function by exploiting weak labels, but other considerations are worth exploring, *e.g.* semi-supervised learning with a few images providing accurate object locations.
- Even if our model provides a philosophy of scale and translation invariant region matching, we can also consider arbitrary rotation by sophisticated region representations (Wang & Kong, 2012).
- Through discussing the parameter λ_2, we interestingly find that some meaningful patches are selected instead of the whole foreground objects. This motivate us to think about learning the discriminative image patches among images similar to (Singh et al., 2012; Doersch et al., 2013). In particular, our work can also benefit part-based model for specific problems, such as fine-grained recognition (Farrell et al., 2011).
- The submodular function also provides a roundabout way for representation learning and instance selection (Krause & Cevher, 2010; Kong & Wang, 2013). Modeling problems with proper submodular function is significantly efficient to deal with large-scale data in practice.

References

Belkin, Mikhail and Niyogi, Partha. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

Bentley, Jon Louis. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18:509–517, 1975.

Boiman, Oren, Shechtman, Eli, and Irani, Michal. In defense of nearest-neighbor based image classification. In *CVPR*, 2008.

Boureau, Y-L, Bach, Francis, LeCun, Yann, and Ponce, Jean. Learning mid-level features for recognition. In *CVPR*, 2010.

Chang, Kai-Yueh, Liu, Tyng-Luh, and Lai, Shang-Hong. From co-saliency to co-segmentation: An efficient and fully unsupervised energy minimization model. In *CVPR*, 2011.

Chen, Jie, Fang, Haw-ren, and Saad, Yousef. Fast approximate k nn graph construction for high dimensional data via recursive lanczos bisection. *JMLR*, 10:1989–2012, 2009.

Coates, Adam and Ng, Andrew. The importance of encoding versus training with sparse coding and vector quantization. In *ICML*, 2011.

Cormen, Thomas H., Leiserson, Charles E., Rivest, Ronald L., and Stein, Clifford. *Introduction to Algorithms*. The MIT Press, 3 edition, 2009.

Dalal, Navneet and Triggs, Bill. Histograms of oriented gradients for human detection. In *CVPR*, 2005.

Dietterich, Thomas G, Lathrop, Richard H, and Lozano-Pérez, Tomás. Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence*, 89 (1):31–71, 1997.

Doersch, Carl, Gupta, Abhinav, and Efros, Alexei A. Mid-level visual element discovery as discriminative mode seeking. In *NIPS*, 2013.

Domingos, Pedro. A few useful things to know about machine learning. *Commun. ACM*, 55(10):78–87, 2012.

Duchenne, Olivier, Joulin, Armand, and Ponce, Jean. A graph-matching kernel for object categorization. In *ICCV*, 2011.

Farrell, Ryan, Oza, Om, Zhang, Ning, Morariu, Vlad I, Darrell, Trevor, and Davis, Larry S. Birdlets: Subordinate categorization using volumetric primitives and pose-normalized appearance. In *ICCV*, 2011.

Fei-Fei, Li, Fergus, Rob, and Perona, Pietro. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *CVIU*, 2007.

Fu, Huazhu, Cao, Xiaochun, and Tu, Zhuowen. Cluster-based co-saliency detection. *TIP*, 22(10):3766–3778, 2013.

Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. *arXiv*, 2013.

Griffin, G., Holub, A., and Perona, P. Caltech-256 Object Category Dataset. Technical Report CNS-TR-2007-001, California Institute of Technology, 2007.

Iyer, Rishabh, Jegelka, Stefanie, and Jeff, Bilmes. fast semidifferential-based submodular function optimization. In *ICML*, 2013.

Jia, Yangqing, Huang, Chang, and Darrell, Trevor. Beyond spatial pyramids: Receptive field learning for pooled image features. In *CVPR*, 2012.

Jiang, Zhuolin, Lin, Zhe, and Davis, Larry S. Label consistent k-svd: Learning a discriminative dictionary for recognition. *PAMI*, 35(11):2651–2664, 2013.

Kim, Gunhee and Xing, Eric P. Jointly aligning and segmenting multiple web photo streams for the inference of collective photo storylines. In *CVPR*, 2013.

Kong, Shu and Wang, Donghui. Learning exemplar-represented manifolds in latent space for classification. In *ECML-PKDD*. 2013.

Kong, Shu, Jiang, Zhuolin, and Yang, Qiang. Learning mid-level features and modeling neuron selectivity for image classification. *arXiv preprint arXiv:1401.5535*, 2014.

Krause, Andreas and Cevher, Volkan. Submodular dictionary selection for sparse representation. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 567–574, 2010.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoff. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

Lazebnik, Svetlana, Schmid, Cordelia, and Ponce, Jean. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, 2006.

Lee, Honglak, Grosse, Roger, Ranganath, Rajesh, and Ng, Andrew Y. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML*, 2009.

Leskovec, Jure, Krause, Andreas, Guestrin, Carlos, Faloutsos, Christos, VanBriesen, Jeanne, and Glance, Natalie. Cost-effective outbreak detection in networks. In *KDD*, 2007.

Lowe, David G. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.

Nemhauser, George L, Wolsey, Laurence A, and Fisher, Marshall L. An analysis of approximations for maximizing submodular set functionsłi. *Mathematical Programming*, 14(1):265–294, 1978.

Nguyen, Minh Hoai, Torresani, Lorenzo, de la Torre, Fernando, and Rother, Carsten. Weakly supervised discriminative localization and classification: a joint learning process. In *ICCV*, 2009.

Olshausen, Bruno A et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 1996.

Russakovsky, Olga, Lin, Yuanqing, Yu, Kai, and Fei-Fei, Li. Object-centric spatial pooling for image classification. In *ECCV*. 2012.

Shi, Jianbo and Malik, Jitendra. Normalized cuts and image segmentation. *PAMI*, 22(8):888–905, 2000.

Singh, Saurabh, Gupta, Abhinav, and Efros, Alexei A. Unsupervised discovery of mid-level discriminative patches. In *ECCV*. 2012.

Tatler, Benjamin W. The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision*, 7(14), 2007.

van de Sande, Koen EA, Uijlings, Jasper RR, Gevers, Theo, and Smeulders, Arnold WM. Segmentation as selective search for object recognition. In *ICCV*, 2011.

Wang, Donghui and Kong, Shu. Learning class-specific dictionaries for digit recognition from spherical surface of a 3d ball. *Machine Vision and Applications*, pp. 1–15, 2012.

Wang, Jinjun, Yang, Jianchao, Yu, Kai, Lv, Fengjun, Huang, Thomas, and Gong, Yihong. Locality-constrained linear coding for image classification. In *CVPR*, 2010.

Yang, Jianchao, Yu, Kai, Gong, Yihong, and Huang, Thomas. Linear spatial pyramid matching using sparse coding for image classification. In *CVPR*, 2009.

Zeiler, Matthew D. and Fergus, Rob. Visualizing and understanding convolutional networks. *arXiv*, 2013.

Zeiler, Matthew D, Taylor, Graham W, and Fergus, Rob. Adaptive deconvolutional networks for mid and high level feature learning. In *ICCV*, 2011.

Appendix: Proof of Proposition 1 and Lemma 1

We rewrite the function $\mathcal{H}(A)$ proposition and lemma as below for presentational convenience:

$$\mathcal{H}(A) = \log\big(\mu + h(\mathbf{S}_{A,A}) - \alpha h(\mathbf{S}_{A,\bar{A}}) - \beta h(\mathbf{S}_{\bar{A},\bar{A}})\big). \quad (10)$$

Proposition: There exist proper α and β that make the proposed function $\mathcal{H}(A)$: $2^{\mathcal{V}} \rightarrow \mathbb{R}$ a monotonically increasing and submodular function.

Lemma: For any $\tau \in \mathbb{R}^*$, $\mathcal{H}(A)$ is a monotonically increasing and submodular function by setting $\alpha = \tau - 1$ and $\beta = \tau$.

Proof: To prove the above proposition and lemma, we just need to prove the lemma, as this lemma simply leads to a possible choice of parameter α and β, both of which can be tuned by the positive scalar $\tau \in \mathbb{R}^*$.

We present an auxiliary function as below:

$$\mathcal{R}(A) = \log\big(\mu + h(\mathbf{S}_{A,\mathcal{V}}) - \tau h(\mathbf{S}_{\bar{A},\mathcal{V}})\big). \quad (11)$$

With the definition over matrix $\mathbf{S}_{A,B} \in \mathbb{R}^{|A|\times|B|}$:

$$h(\mathbf{S}_{A,B}) = \begin{cases} \sum_{i\in A}\sum_{j\in B} c_{ij}, & A \neq \varnothing \text{ and } B \neq \varnothing, \\ 0, & otherwise. \end{cases} \quad (12)$$

we can easily derive that:

$$\begin{aligned}\mathcal{R}(A) &= \log\big(\mu + h(\mathbf{S}_{A,\mathcal{V}}) - \tau h(\mathbf{S}_{\bar{A},\mathcal{V}})\big) \\ &= \log\big(\mu + h(\mathbf{S}_{A,A}) + h(\mathbf{S}_{A,\bar{A}}) - \tau h(\mathbf{S}_{\bar{A},A}) - \tau h(\mathbf{S}_{\bar{A},\bar{A}})\big) \\ &= \log\big(\mu + h(\mathbf{S}_{A,A}) - (\tau-1) h(\mathbf{S}_{A,\bar{A}}) - \tau h(\mathbf{S}_{\bar{A},\bar{A}})\big).\end{aligned} \quad (13)$$

Therefore, with this auxiliary function, we have $\alpha = \tau - 1$ and $\beta = \tau$ that satisfy $\mathcal{H}(A)$. Now, to prove the proposition and the lemma, we can turn to show $\mathcal{G}(A)$ is a monotonically increasing and submodular function.

To this end, we split the parts in $\mathcal{R}(A)$ as below:

$$\begin{aligned}\mathcal{R}(A) &= \log\big(\mu + h(\mathbf{S}_{A,\mathcal{V}}) - \tau h(\mathbf{S}_{\bar{A},\mathcal{V}})\big) \\ &= \log\big(\mu + \sum_{\substack{i\in A\\ j\in\mathcal{V}}} s_{ij} - \tau \sum_{\substack{i\in \bar{A}\\ j\in\mathcal{V}}} s_{ij}\big) \\ &= \log\big(\mu\mu + \sum_{\substack{i\in A\\ j\in\mathcal{V}}} s_{ij} - \tau \sum_{\substack{i\in \bar{A}\\ j\in\mathcal{V}}} s_{ij} - \tau \sum_{\substack{i\in A\\ j\in\mathcal{V}}} s_{ij} + \tau \sum_{\substack{i\in A\\ j\in\mathcal{V}}} s_{ij}\big) \\ &= \log\big(\mu + (\tau+1) \sum_{\substack{i\in A\\ j\in\mathcal{V}}} s_{ij} - \tau \sum_{\substack{i\in \mathcal{V}\\ j\in\mathcal{V}}} s_{ij}\big) \\ &= \log\big(\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})\big).\end{aligned} \quad (14)$$

monotonically increasing: By definition, if $A = \varnothing$, then we have $\mathcal{R}(A) \geq 0$; moreover, by denoting $\Delta = \mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})$, we have:

$$\begin{aligned}&\mathcal{R}(A \cup a) - \mathcal{R}(A) \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A\cup a,\mathcal{V}})}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})} \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1)\big(h(\mathbf{S}_{A,\mathcal{V}}) + \sum_j s_{aj}\big)}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})} \\ =& \log \frac{\Delta + (\tau+1)\sum_j s_{aj}}{\Delta} \\ =& \log\big(1 + \frac{(\tau+1)\sum_j s_{aj}}{\Delta}\big) \\ \geq& 0.\end{aligned} \quad (15)$$

Therefore, $\mathcal{R}(A)$ is monotonically increasing.

Submodulary: Previous derivation leads to the following:

$$\begin{aligned}&\mathcal{R}(A \cup a) - \mathcal{R}(A) \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A\cup a,\mathcal{V}})}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})} \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1)\big(h(\mathbf{S}_{A,\mathcal{V}}) + \sum_j s_{aj}\big)}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})} \\ =& \log \frac{\Delta + (\tau+1)\sum_j s_{aj}}{\Delta}.\end{aligned} \quad (16)$$

Similarly, we derive:

$$\begin{aligned}&\mathcal{R}(A \cup a \cup b) - \mathcal{R}(A \cup b) \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A\cup a\cup b,\mathcal{V}})}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A\cup b,\mathcal{V}})} \\ =& \log \frac{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1)\big(h(\mathbf{S}_{A,\mathcal{V}}) + \sum_j s_{aj} + \sum_j s_{bj}\big)}{\mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1)\big(h(\mathbf{S}_{A,\mathcal{V}}) + \sum_j s_{bj}\big)} \\ =& \log \frac{\Delta + (\tau+1)(\sum_j s_{aj} + \sum_j s_{bj})}{\Delta + (\tau+1)\sum_j s_{bj}}\end{aligned} \quad (17)$$

where $\Delta = \mu - \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}}) + (\tau+1) h(\mathbf{S}_{A,\mathcal{V}})$. Let $\mu = 1 + \tau h(\mathbf{S}_{\mathcal{V},\mathcal{V}})$, and denote $x_1 = \Delta$ and $x_2 = \Delta + (\tau +$

$1)\sum_j s_{bj}$, we have $0 \leq x_1 \leq x_2$, and thus:

$$\begin{aligned}
&x_1 \leq x_2 \\
\Longleftrightarrow &\frac{1}{x_1} \geq \frac{1}{x_2} \\
\Longleftrightarrow &\frac{(\tau+1)\sum_j s_{aj}}{x_1} \geq \frac{(\tau+1)\sum_j s_{aj}}{x_2} \\
\Longleftrightarrow &1 + \frac{(\tau+1)\sum_j s_{aj}}{x_1} \geq 1 + \frac{(\tau+1)\sum_j s_{aj}}{x_2} \\
\Longleftrightarrow &\frac{x_1 + (\tau+1)\sum_j s_{aj}}{x_1} \geq \frac{x_2 + (\tau+1)\sum_j s_{aj}}{x_2} \\
\Longleftrightarrow &\log\frac{x_1 + (\tau+1)\sum_j s_{aj}}{x_1} \geq \log\frac{x_2 + (\tau+1)\sum_j s_{aj}}{x_2} \\
\Longleftrightarrow &\big(\mathcal{G}(A\cup a) - \mathcal{G}(A)\big) \geq \big(\mathcal{G}(A\cup a\cup b) - \mathcal{G}(A\cup b)\big).
\end{aligned} \tag{18}$$

Therefore, the auxiliary function $\mathcal{R}(A)$ is submodular.

Summary: Since the auxiliary $\mathcal{R}(A)$ is a monotonically increasing and submodular function, with the relationship between τ and $\alpha = \tau - 1$ and $\beta = \tau$, we show the derived $\mathcal{H}_\tau(A)$ is a monotonically increasing and submodular function.

End of proof. ■

Appendix: Proof of Proposition 2

Proposition: With the function defined as below over vector $\mathbf{x} = [x_1, \ldots, x_N] \succcurlyeq 0$):

$$g(\mathbf{x}) = \sum_{j=1}^{N} \log(x_j + 1),$$

if $x_c \leq x_i, \forall i \neq c$, then $g(x_c + 1|\mathbf{x}) \geq g(x_i + 1|\mathbf{x})$, where $g(x_i + 1|\mathbf{x})$ is defined as: $g(x_i + 1|\mathbf{x}) = \sum_{j\neq i} \log(x_j + 1) + \log(x_i + 1 + 1)$.

Proof: With $0 \leq x_c \leq x_i$, $\forall i \neq c$, we writing down $g(x_c + 1|\mathbf{x})$ and $g(x_i + 1|\mathbf{x})$ as below:

$$\begin{aligned}
g(x_c + 1|\mathbf{x}) &= \sum_{j\neq c, j\neq i} \log(x_j + 1) + \log(x_c + 1 + 1) + \log(x_i + 1), \\
g(x_i + 1|\mathbf{x}) &= \sum_{j\neq c, j\neq i} \log(x_j + 1) + \log(x_c + 1) + \log(x_i + 1 + 1).
\end{aligned} \tag{19}$$

Then, we have:

$$\begin{aligned}
&g(x_c + 1|\mathbf{x}) - g(x_i + 1|\mathbf{x}) \\
=&\log(x_c + 1 + 1) + \log(x_i + 1) - \log(x_c + 1) - \log(x_i + 1 + 1) \\
=&\log(x_c + 1 + 1)(x_i + 1) - \log(x_c + 1)(x_i + 1 + 1) \\
=&\log\frac{(x_c + 1 + 1)(x_i + 1)}{(x_c + 1)(x_i + 1 + 1)} \\
=&\log\frac{(x_c + 1)(x_i + 1) + (x_i + 1)}{(x_c + 1)(x_i + 1) + (x_c + 1)}.
\end{aligned} \tag{20}$$

As $0 \leq x_c \leq x_i, \forall i \neq c$, we have:

$$\begin{aligned}
&x_c \leq x_i \\
\Longleftrightarrow &(x_c + 1)(x_i + 1) + 1 + x_c \leq (x_c + 1)(x_i + 1) + 1 + x_i \\
\Longleftrightarrow &\frac{(x_c + 1)(x_i + 1) + (x_i + 1)}{(x_c + 1)(x_i + 1) + (x_c + 1)} \geq 1 \\
\Longleftrightarrow &\log\frac{(x_c + 1)(x_i + 1) + (x_i + 1)}{(x_c + 1)(x_i + 1) + (x_c + 1)} \geq 0 \\
\Longleftrightarrow &g(x_c + 1|\mathbf{x}) - g(x_i + 1|\mathbf{x}) \geq 0.
\end{aligned} \tag{21}$$

End of proof. ■

Appendix: Proof of Proposition 3

The proposed function as below to maximize is a monotonically increasing and submodular function:

$$\begin{aligned}
\mathcal{G}(A) = &\sum_{j=1}^{N} \log(|A_j| + 1), \\
\text{s.t. } &\cup_{j=1}^{N} A_j = A, A_i \cap A_j = \varnothing, \forall i \neq j.
\end{aligned} \tag{22}$$

Proof: To prove $\mathcal{G}(A)$ is monotonically increasing, we just need to show $\mathcal{G}(A \cup a) - \mathcal{G}(A) \geq 0$, where $a \in \mathcal{V}$ and $a \notin A$. Without of lose of generality, we can suppose a comes from the i^{th} image (note $a \notin A_i$), therefore a is added to A_i. With simple derivations, we have:

$$\mathcal{G}(A \cup a) - \mathcal{G}(A) = \log(|A_i + a| + 1) - \log(|A_i| + 1) \tag{23}$$

Here $|\cdot|$ means the cardinality. we can denote $x = |A_i| + 1$ is a positive integer, hence it is easy to see:

$$\mathcal{G}(A \cup a) - \mathcal{G}(A) = \log(x + 1) - \log(x) > 0. \tag{24}$$

Therefore, $\mathcal{G}(A)$ is a strictly monotonically increasing function.

Moreover, for its submodularity, we need to show the following for $a \notin A$ and $b \notin A$ ($a \neq b$, otherwise equality is achieved):

$$\mathcal{G}(A \cup a) - \mathcal{G}(A) \geq \mathcal{G}(A \cup b \cup a) - \mathcal{G}(A \cup b). \tag{25}$$

There are two cases, a and b come from a common image, or two different ones.

For the first case, suppose a and b come from the i^{th} image,

then we have:

$$\begin{aligned}
&\big(\mathcal{G}(A\cup a)-\mathcal{G}(A)\big)-\big(\mathcal{G}(A\cup b\cup a)-\mathcal{G}(A\cup b)\big)\\
=&\mathcal{G}(A\cup a)-\mathcal{G}(A)-\mathcal{G}(A\cup b\cup a)+\mathcal{G}(A\cup b)\\
=&\sum_{j\neq i}^{N}\log(|A_j|+1)+\log(|A_i\cup a|+1)\\
&-\sum_{j\neq i}^{N}\log(|A_j|+1)-\log(|A_i|+1)\\
&-\sum_{j\neq i}^{N}\log(|A_j|+1)-\log(|A_i\cup a\cup b|+1)\\
&+\sum_{j\neq i}^{N}\log(|A_j|+1)+\log(|A_i\cup b|+1)\\
=&\log(|A_i\cup a|+1)-\log(|A_i|+1)\\
&-\log(|A_i\cup a\cup b|+1)+\log(|A_i\cup b|+1)\\
=&\big(\log(x+1)-\log(x)\big)-\big(\log(x+2)-\log(x+1)\big),
\end{aligned}\tag{26}$$

where x is a positive integer. Now, the question turns to proving $\log(x+1)$ is a concave function. This is obvious, and thus proof done.

For the second case, *i.e.* a and b come from different images, we assume a and b come from the i^{th} and the k^{th} image, respectively. Then we have:

$$\begin{aligned}
&(\mathcal{G}(A\cup a)-\mathcal{G}(A))-(\mathcal{G}(A\cup b\cup a)-\mathcal{G}(A\cup b))\\
=&\mathcal{G}(A\cup a)-\mathcal{G}(A)-\mathcal{G}(A\cup b\cup a)+\mathcal{G}(A\cup b)\\
=&\sum_{j\neq i}^{N}\log(|A_j|+1)+\log(|A_i\cup a|+1)\\
&-\sum_{j\neq i}^{N}\log(|A_j|+1)-\log(|A_i|+1)\\
&-\sum_{j\neq i,j\neq k}^{N}\log(|A_j|+1)-\log(|A_i\cup a|+1)-\log(|A_k\cup b|+1)\\
&+\sum_{j\neq i,j\neq k}^{N}\log(|A_j|+1)+\log(|A_i|+1)+\log(|A_k\cup b|+1)\\
=&\log(|A_i\cup a|+1)-\log(|A_i|+1)\\
&-\log(|A_i\cup a|+1)+\log(|A_i|+1)\\
=&0
\end{aligned}\tag{27}$$

Therefore, in this case, $\mathcal{G}(A)$ is a modular function.

Overall, we prove the function $\mathcal{G}(A)$ is a monotonically increasing and submodular function.

End of proof.■

Appendix: Proof of Proposition 4

We rewrite the function $\mathcal{F}(A)$ and the proposition as below for presentational convenience:

$$\mathcal{F}(A)\equiv\mathcal{H}_\tau(A)+\lambda_1\mathcal{G}(A)+\lambda_2\|\mathbf{q}_{(A)}\|_1. \tag{28}$$

Proposition: The proposed function $\mathcal{F}(A)$ is a monotonically increasing and submodular function, and induces a matroid $\mathcal{M}=(\mathcal{V},\mathcal{I})$, where $\mathcal{V}$ is the point set, and $\mathcal{I}$ is the collection of subsets $A\subseteq\mathcal{V}$.

Proof: As we previously show $\mathcal{H}_\tau(A)$ and $g(A)$ are monotonically increasing and submodular functions, now we just need to fucus on $\|\mathbf{q}_A\|_1$. By definition, $q_i\geq 0,\forall i$, and

$$\|\mathbf{q}_A\|_1=\sum_{i\in A}q_i. \tag{29}$$

It is easy to see:

$$\|\mathbf{q}_{A\cup a}\|_1-\|\mathbf{q}_A\|_1=q_a\geq 0, \tag{30}$$

thus $\|\mathbf{q}_A\|_1$ is monotonically increasing.

Furthermore, we have:

$$\begin{aligned}
&\big(\|\mathbf{q}_{A\cup a}\|_1-\|\mathbf{q}_A\|_1\big)-\big(\|\mathbf{q}_{A\cup a\cup b}\|_1-\|\mathbf{q}_{A\cup b}\|_1\big)\\
=&q_a-q_a\\
=&0.
\end{aligned}\tag{31}$$

Therefore, $\|\mathbf{q}_A\|_1$ is a modular function.

In sum, $\|\mathbf{q}_A\|_1$ is a monotonically increasing and modular function; and thus, the objective function is a monotonically increasing and submodular function.

matroid: The proposed objective function induces a matroid $\mathcal{M}=(\mathcal{V},\mathcal{I})$, where $\mathcal{V}$ is the ground set, and $\mathcal{I}$ is a family of feasible solution sets.

Proof focuses on the following three conditions:

1. $\varnothing\in\mathcal{I}$: the function start with $\varnothing$ as defined.
2. (Hereditary property): If $A\subseteq B$ and $B\in\mathcal{I}$, then $A\in\mathcal{I}$;
3. (Exchange property): If $A\in\mathcal{I}, B\in\mathcal{I}$ and $|A|<|B|$, there is an element $e\in B-A$ such that $A\cup e\in\mathcal{I}$.

As there is no constraint on the matroid posed in $\mathcal{F}(A)$, our objective function induces the desired set A from a uniform matroid.

End of proof. ■

Appendix: Other Details in Implementation

A trick to build the dissimilarity graph is to correlate each pair of images with their RF's by selecting fixed number (say 3) of nearest RF's with brutal force, then we smooth the graph by keeping the fixed number (say the number of training images of each category) of entries with the smallest values and derive the kNN graph.

Appendix: More Results on the Synthetic Data

Fig. 7 presents more illustrations on the marginal gains of different iterations. We can see after sufficient iterations, all the RF_+'s that are assumed to be most correlated are found finally.

Appendix: More Results of Caltech256

Fig. 8 displays more results of the learned receptive fields over images from Caltech256. From the figure, we can see the most informative RF's are found in the images. This demonstrates the effectiveness of our method.

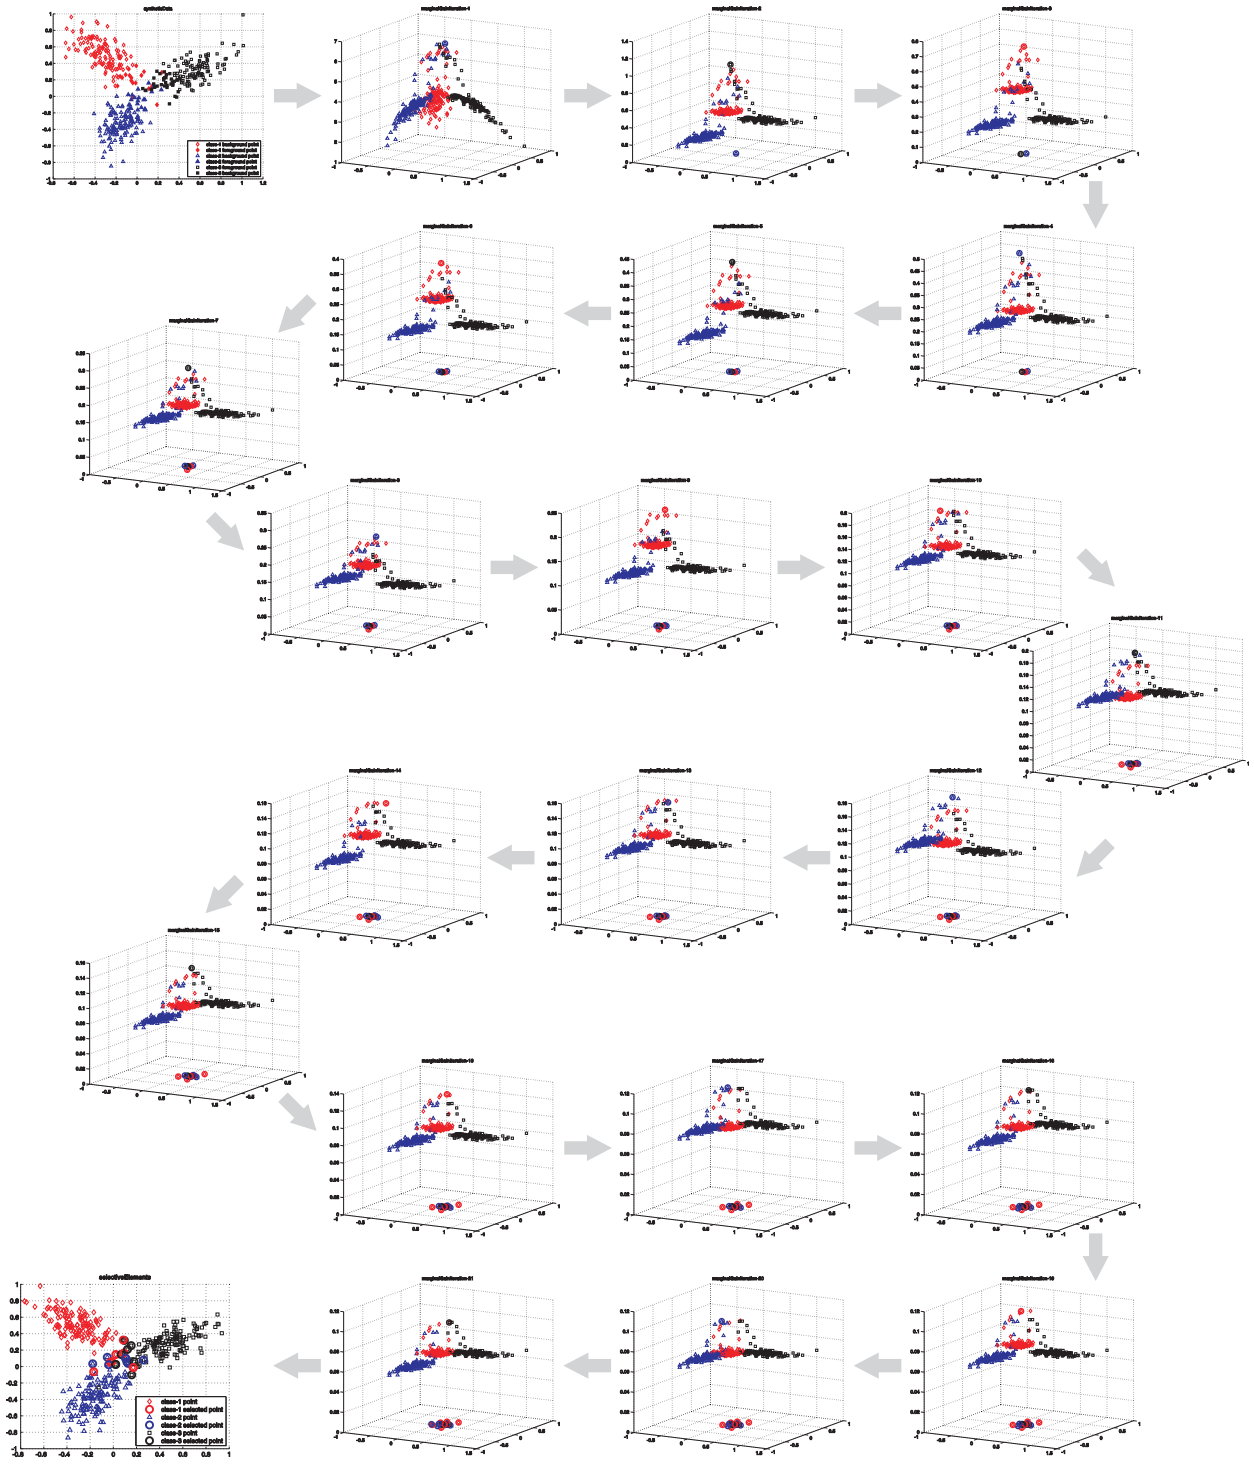

Figure 7. More iterations over the synthetic dataset to select more receptive fields. It can be seen that all the selected data lie on the intersection of the three classes/clusters. This demonstrates the effectiveness of the proposed method.

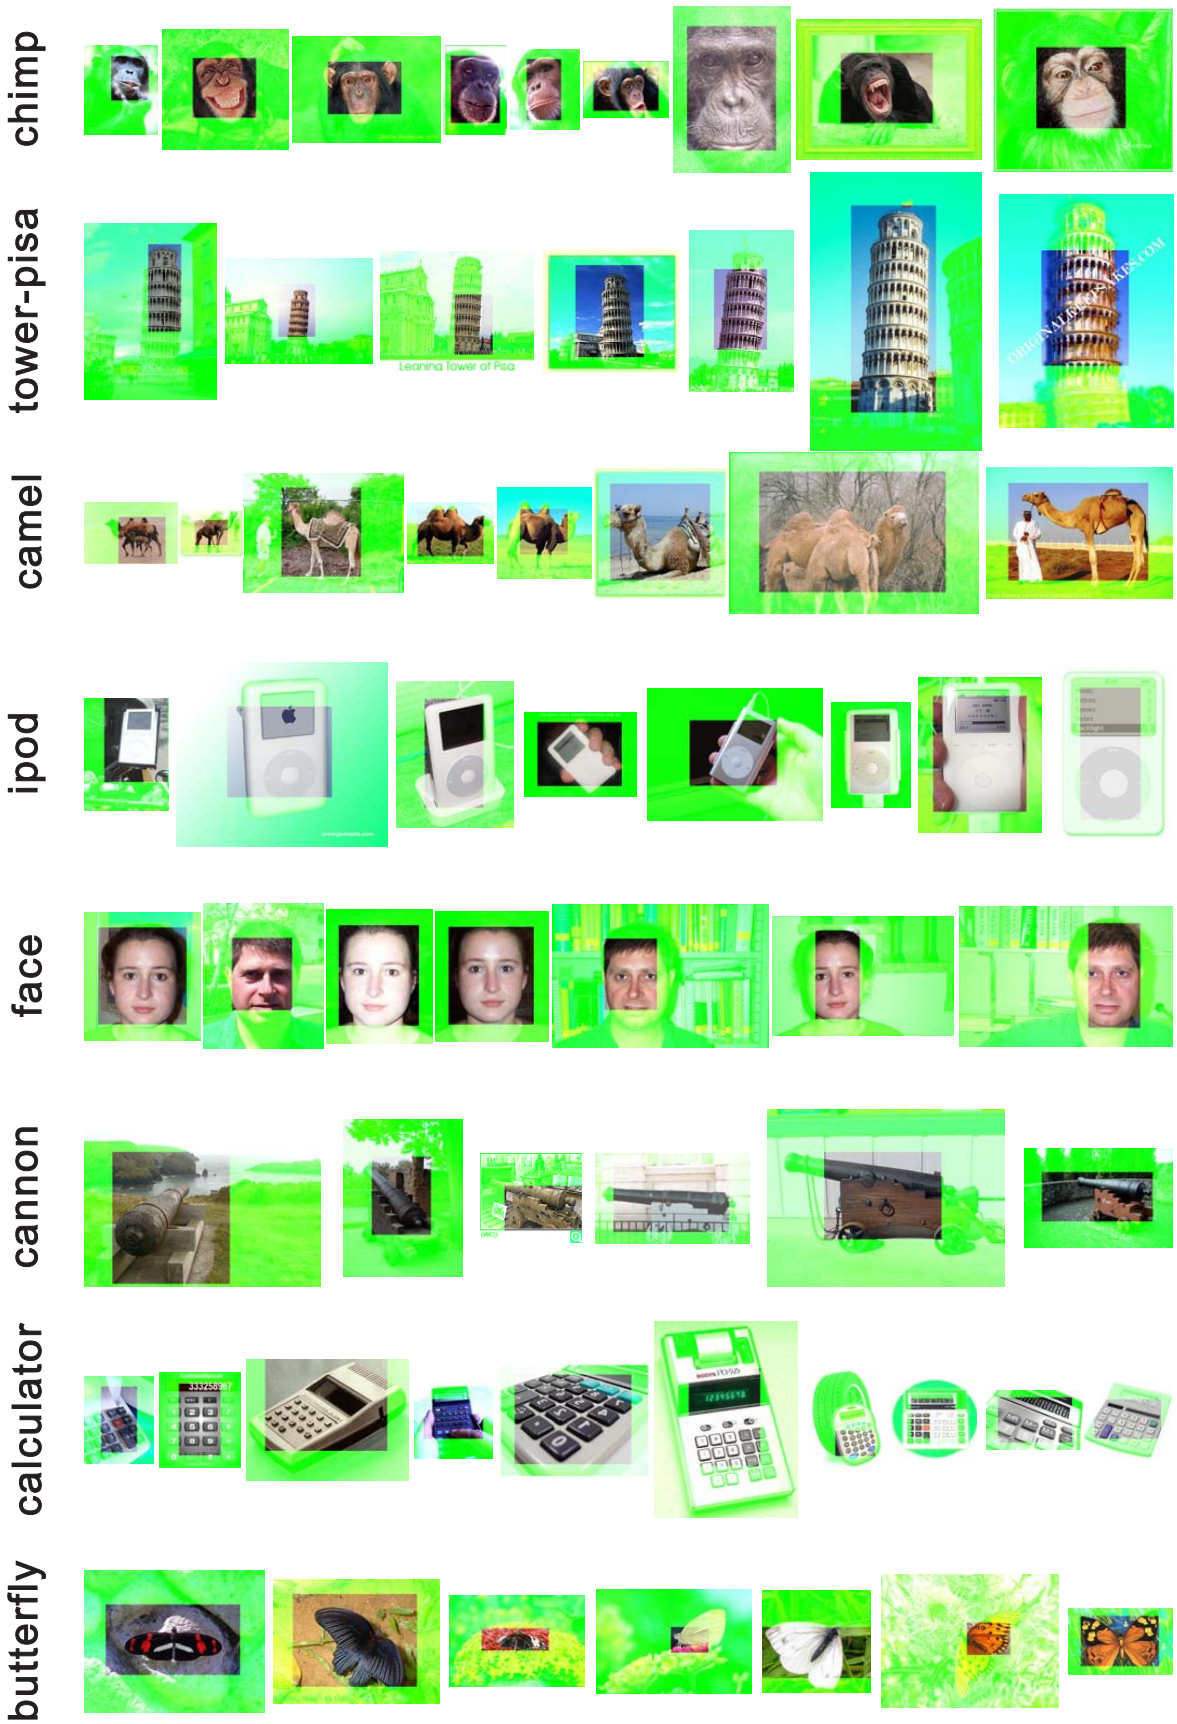

Figure 8. (Best seen in color.) Further illustration of the learned receptive fields over images from Caltech256.